

Richard C. Carrier, Ph.D.

“Bayes’ Theorem for Beginners: Formal Logic and Its Relevance to Historical Method — Adjunct Materials and Tutorial”

The Jesus Project Inaugural Conference
“Sources of the Jesus Tradition: An Inquiry”

5 December 2008 (Amherst, NY)

Table of Contents for Enclosed Document

Handout Accompanying Oral Presentation of December 5.....pp. 2-5

Adjunct Document Expanding on Oral Presentation.....pp. 6-26

Simple Tutorial in Bayes’ Theorem.....pp. 27-39

NOTE: A chapter of the same title will be published by Prometheus Press (in *Sources of the Jesus Tradition*) discussing or referring to the contents of this online document. That primary document (to which this document is adjunct) has also been published in advance as **“Bayes’ Theorem for Beginners: Formal Logic and Its Relevance to Historical Method”** in *Caesar: A Journal for the Critical Study of Religion and Human Values* 3.1 (2009): 26-35.

Richard C. Carrier, Ph.D.

"Bayes Theorem for Beginners: Formal Logic and Its Relevance to Historical Method"

December 2008 (Amherst, NY)

Notes and Bibliography

1. Essential Reading on "Historicity Criteria"

Stanley Porter, *The Criteria for Authenticity in Historical-Jesus Research: Previous Discussion and New Proposals* (Sheffield Academic Press: 2000).

Christopher Tuckett, "Sources and Methods," *The Cambridge Companion to Jesus*, edited by Markus Bockmuehl (Cambridge University Press: 2001): pp. 121-37.

Gerd Theissen and Dagmar Winter, *The Quest for the Plausible Jesus: The Question of Criteria* (John Knox Press: 2002).

2. Example List of Popular Historicity Criteria

Incomplete List (names often differ, criteria often overlap – here are 17; there are two or three dozen):

Dissimilarity	- dissimilar to independent Jewish or Christian precedent
Embarrassment	- if it was embarrassing, it must be true
Coherence	- coheres with other confirmed data
Multiple Attestation	- attested in more than one independent source
Contextual Plausibility	- plausible in a Jewish or Greco-Roman cultural context
Historical Plausibility	- coheres with a plausible historical reconstruction
Natural Probability	- coheres with natural science (etc.)
Explanatory Credibility	- historicity better explains later traditions
Oral Preservability	- capable of surviving oral transmission
Fabricatory Trend	- isn't part of known trends in fabrication or embellishment
Least Distinctiveness	- the simpler version is the more historical
Vividness of Narration	- the more vivid, the more historical
Crucifixion	- explains why Jesus was crucified
Greek Context	- if whole context suggests parties speaking Greek
Aramaic Context	- if whole context suggests parties speaking Aramaic
Textual Variance	- the more invariable a tradition, the more historical
Discourse Features	- if J's speeches cohere in style but differ fr. surrounding text

3. Formal Logical Analysis – discussion & examples available by online document:

www.richardcarrier.info/CarrierDec08.pdf

4. Recommended Reading on Bayes' Theorem

Eliezer Yudkowsky, "An Intuitive Explanation of Bayesian Reasoning (Bayes' Theorem for the Curious and Bewildered: An Excruciatingly Gentle Introduction)" at yudkowsky.net/bayes/bayes.html

Douglas Hunter, *Political [and] Military Applications of Bayesian Analysis: Methodological Issues* (Westview Press: 1984).

Luc Bovens and Stephan Hartmann, *Bayesian Epistemology* (Oxford: 2003).

Timothy McGrew, "Bayesian Reasoning: An Annotated Bibliography" at homepages.wmich.edu/~mcgrew/bayes.htm

Wikipedia on "Bayes' Theorem" (for English: en.wikipedia.org/wiki/Bayes'_theorem)

5. Bayes' Theorem (Complete)

$$P(h|e.b) = \frac{P(h|b) \times P(e|h.b)}{[P(h|b) \times P(e|h.b)] + [P(\sim h|b) \times P(e|\sim h.b)]}$$

6. Bayes' Theorem (Abbreviated)

$$P(h|e.b) = \frac{P(h|b) \times P(e|h.b)}{P(e|b)}$$

Note that $P(e|b) = [P(h|b) \times P(e|h.b)] + [P(\sim h|b) \times P(e|\sim h.b)]$ so I recommend the longer form instead, as a check against error (especially that of assuming $P(e|b) = P(e|\sim h.b)$).

7. Explanation of the Terms in Bayes' Theorem

P = Probability (epistemic probability = the probability that something stated is true)

h = hypothesis being tested

~h = all other hypotheses that could explain the same evidence (if *h* is false)

e = all the evidence directly relevant to the truth of *h* (*e* includes both what *is* observed *and* what is *not* observed)

b = total background knowledge (all available personal and human knowledge about anything and everything, from physics to history)

$P(h|e.b)$ = the probability that a hypothesis (*h*) is true given all the available evidence (*e*) and all our background knowledge (*b*)

$P(h|b)$ = **the prior probability that h is true** = the probability that our hypothesis would be true given only our background knowledge (i.e. if we knew nothing about e)

$P(e|h.b)$ = **the posterior probability of the evidence (given h and b)** = the probability that all the evidence we have would exist (or something comparable to it would exist) if the hypothesis (and background knowledge) is true. = [consequent probability]

$P(\sim h|b)$ = $1 - P(h|b)$ = **the prior probability that h is false** = the sum of the prior probabilities of all alternative explanations of the same evidence (e.g. if there is only one viable alternative, this means the prior probability of all other theories is vanishingly small, i.e. substantially less than 1%, so that $P(\sim h|b)$ is the prior probability of the one viable competing hypothesis; if there are many viable competing hypotheses, they can be subsumed under one group category ($\sim h$), or treated independently by expanding the equation, e.g. for three competing hypotheses [$P(h|b) \times P(e|h.b)$] + [$P(\sim h|b) \times P(e|\sim h.b)$] becomes [$P(h_1|b) \times P(e|h_1.b)$] + [$P(h_2|b) \times P(e|h_2.b)$] + [$P(h_3|b) \times P(e|h_3.b)$])

$P(e|\sim h.b)$ = **the posterior probability of the evidence if b is true but h is false** = the probability that all the evidence we have would exist (or something comparable to it would exist) if the hypothesis we are testing *is false*, but all our background knowledge is still true. This also equals the posterior probability of the evidence if some hypothesis *other* than h is true—and if there is more than one viable contender, you can include each competing hypothesis independently (per above) or subsume them all under one group category ($\sim h$). = [consequent probability]

Since mathematicians use posterior probability to refer to the final probability (the probability of h posterior to e), to avoid confusion I will hereafter use consequent probability to refer to what I here describe as the posterior probability of the evidence (the probability of e posterior to h).

8. The Advantages of Bayes' Theorem

1. Helps to tell if your theory is *probably* true rather than merely *possibly* true
2. Inspires closer examination of your backgr. knowl. and assumptions of likelihood
3. Forces examination of the likelihood of the evidence on competing theories
4. Eliminates the Fallacy of Diminishing Probabilities
5. Bayes' Theorem has been proven to be formally valid
6. Bayesian reasoning with or without math exposes assumptions to criticism & consequent revision and therefore promotes progress

9. Avoiding Common Errors with Bayes' Theorem

1. The Fallacy of False Precision

(SOLUTION: include reasonable margins of error)

2. The Fallacy of Confusing Evidence with Theories

(SOLUTION: try to limit contents of e to tangible physical facts, i.e. actual surviving artifacts, documents, etc., and straightforward generalizations therefrom)

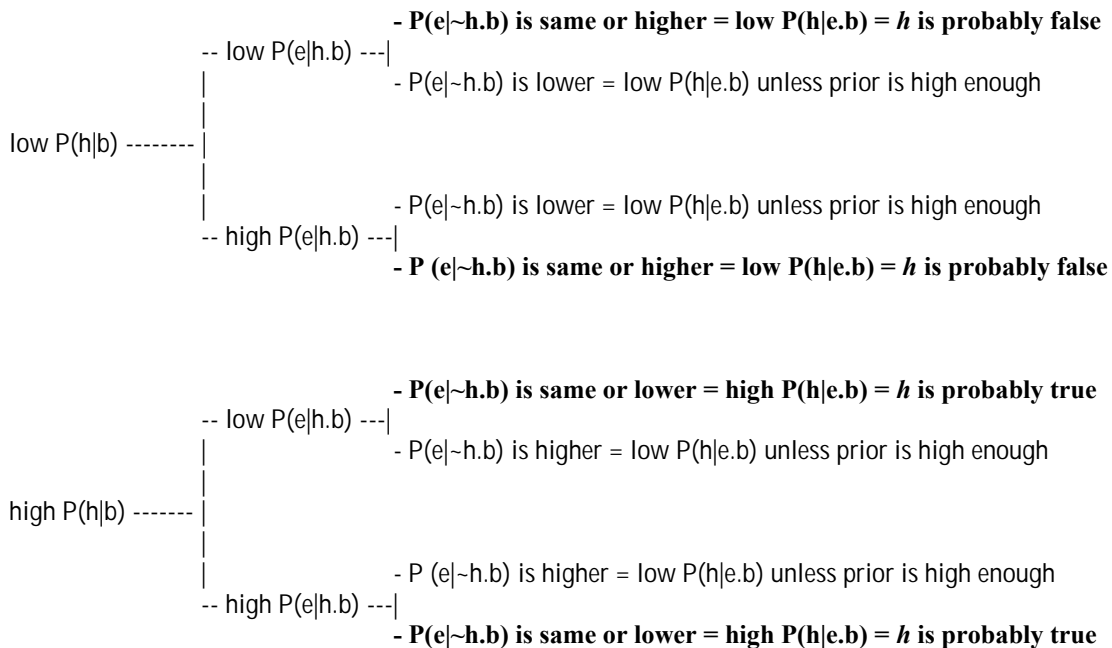
3. The Fallacy of Confusing Assumptions with Knowledge

(SOLUTION: try to limit contents of b to universally agreed expert knowledge)

NOTE: Further discussion of all these topics (and in considerable detail) will appear in my forthcoming book: Richard Carrier, *On the Historicity of Jesus Christ* (there especially in chapter two and a mathematical appendix). For contact information and more about my work in other aspects of Jesus studies and beyond, see www.richardcarrier.info.

To employ Bayesian reasoning without doing any math, employ relative (but non-numerical) estimates of likelihood on the following flowchart:

Bayesian Reasoning Flowchart



"Low" and "High" = lower than 0.5 (50%) and higher than 0.5 (50%), respectively ; when $P(h|b) = 0.5$, so does $P(\sim h|b)$: then the hypothesis with the higher $P(e|b)$ is probably true.

"Prior is high enough" = when $P(h|b)$ is higher than the Bayesian ratio between either $P(e|\sim h.b)$ and $P(e|h.b)$ or vice versa, enough to overcome the gap and thus produce the opposite result.

Richard C. Carrier, Ph.D.

"Bayes Theorem for Beginners: Formal Logic and Its Relevance to Historical Method"

December 2008 (Amherst, NY)

At our Jesus Project conference I will refer to items in this document as I go, in lieu of slides. But I won't discuss every item in it.

1. Essential Reading on "Historicity Criteria"

Stanley Porter, *The Criteria for Authenticity in Historical-Jesus Research: Previous Discussion and New Proposals* (Sheffield Academic Press: 2000).

Christopher Tuckett, "Sources and Methods," *The Cambridge Companion to Jesus*, edited by Markus Bockmuehl (Cambridge University Press: 2001): pp. 121-37.

Gerd Theissen and Dagmar Winter, *The Quest for the Plausible Jesus: The Question of Criteria* (John Knox Press: 2002).

2. Examining Historicity Criteria

Typical Problems:

1. The criterion is invalidly applied (e.g. the text actually fails to fulfill the criterion, contrary to a scholar's assertion or misapprehension)
2. The criterion itself is invalid (e.g. the criterion depends upon a rule of inference that is inherently fallacious, contrary to a scholar's intuition)

Required Solutions:

1. Scholars are obligated to establish with clear certainty that a particular item actually fulfills any stated criterion (and what exactly that criterion is).
2. Scholars are obligated to establish the formal logical validity of any stated criterion (especially if establishing its validity requires adopting for it a set of qualifications or conditions previously overlooked).

Incomplete List (names often differ, criteria often overlap – here are 17; there are two or three dozen):

Dissimilarity	- dissimilar to independent Jewish or Christian precedent
Embarrassment	- if it was embarrassing, it must be true
Coherence	- coheres with other confirmed data
Multiple Attestation	- attested in more than one independent source
Contextual Plausibility	- plausible in a Jewish or Greco-Roman cultural context
Historical Plausibility	- coheres with a plausible historical reconstruction
Natural Probability	- coheres with natural science (etc.)
Explanatory Credibility	- historicity better explains later traditions
Oral Preservability	- capable of surviving oral transmission
Fabricatory Trend	- isn't part of known trends in fabrication or embellishment
Least Distinctiveness	- the simpler version is the more historical
Vividness of Narration	- the more vivid, the more historical
Crucifixion	- explains why Jesus was crucified
Greek Context	- if whole context suggests parties speaking Greek
Aramaic Context	- if whole context suggests parties speaking Aramaic
Textual Variance	- the more invariable a tradition, the more historical
Discourse Features	- if J's speeches cohere in style but differ fr. surrounding text

3. Formal Logic: The Basic Syllogism

Major Premise: [Some general rule]

Minor Premise: [Some specific fact satisfying the general rule]

Conclusion: [That which follows necessarily from the major and minor premise]

Strongest Argument:

Major Premise: All working wagons had wheels.

Minor Premise: Jacob owned a working wagon.

Conclusion: Therefore, Jacob owned a wheel.

[This is true by virtue of the definition of 'working wagon']

Stronger Argument:

Major Premise: All major cities in antiquity had sewers.

Minor Premise: Jerusalem was a major city in antiquity.

Conclusion: Therefore, Jerusalem had sewers.

[This is true to a very high degree of probability by virtue of abundant background evidence, including archaeological and textual, that is uncontested for its scope and degree]

= *Major Premise:* [Very probably] all major cities in antiquity had sewers.

Minor Premise: [Very probably] Jerusalem was a major city in antiquity.

Conclusion: Therefore, [very probably] Jerusalem had sewers.

Weaker Argument:

Major Premise: All major cities in antiquity had public libraries.

Minor Premise: Jerusalem was a major city in antiquity.

Conclusion: Therefore, Jerusalem had a public library.

[This is true but not to a high degree of probability by virtue of archaeological and textual evidence that is contestable to some degree]

= *Major Premise:* [Most] major cities in antiquity had public libraries.
= [Somewhat probably] any major city in antiquity had a public library.
Minor Premise: [Very probably] Jerusalem was a major city in antiquity.
Conclusion: Therefore, [somewhat probably] Jerusalem had a public library.

[Here basic theory would produce a probability for the conclusion, $P(\text{Jerusalem had a public library}) = P(JL)$, equal to the *product* of the probabilities of the major and minor premise, due to the Law of Conditional Probability: so although one might say the conclusion is "[somewhat probably] Jerusalem had a public library" this "somewhat probably" will be slightly less than the "somewhat probably" implied in the major premise. For example, if the major premise has a probability of 60% and the minor premise a probability of 90%, then $P(JL) = 0.6 \times 0.9 = 0.54 = 54\%$, not 60%. This can lead to a fallacy of diminishing probabilities, where the more evidence you have, the lower the probability of the conclusion (see next), which clearly cannot be correct. The solution is Bayes' Theorem, which eliminates the fallacy of diminishing probabilities. Hence the importance of Bayes' Theorem.]

Complex Argument:

Major Premise 1: [Probably] all major cities in antiquity had public libraries.

Major Premise 2: [Probably] when the surviving text of something written by an ancient author mentions consulting books in a city's public library, then that city had a public library.

Minor Premise 1: [Very probably] Jerusalem was a major city in antiquity.

Minor Premise 2: [Very probably] we have a 3rd century papyrus fragment of an encyclopedia written by Hippolytus a few decades before, which mentions his consulting books in Jerusalem's public library.

Conclusion: Therefore, [very probably] Jerusalem had a public library.

[Here the probability that Jerusalem had a public library should be increased by the fact of having two kinds of evidence mutually supporting the same conclusion, including both general evidence, and specific evidence. More evidence of either kind could be added to raise the probability of the conclusion even more. But observe, if $P(\text{Major 1}) = 60\%$ and $P(\text{Major 2}) = 60\%$ and $P(\text{Minor 1}) = 90\%$ and $P(\text{Minor 2}) = 90\%$, we get $P(\text{Conclusion}) = P(\text{Jerusalem had a public library}) = P(JL) = 0.6 \times 0.6 \times 0.9 \times 0.9 = 0.29 = 29\%$. Thus the conclusion appears to be less probable when we get more evidence, which clearly cannot be correct. Therefore we need Bayes' Theorem, which avoids this fallacious application of probabilities.]

4. Nesting as a Method of Verification

It's easy to determine if a syllogism is valid (just observe if the conclusion follows from the premises), but harder to determine if a syllogism is sound (since that requires all the premises to be true—as a conclusion can only be as true as its weakest premise). A good method to check for errors in the way you derive confidence in your premises (and thus in how you determine a premise is true), is to build out the required syllogisms supporting each premise, nesting one set of syllogisms within the other.

For example:

Major Premise 1: All major cities in antiquity had public libraries.

Minor Premise 1: Jerusalem was a major city in antiquity.

Conclusion: Therefore, Jerusalem had a public library.

Nesting Test for Major Premise 1:

Major Premise 2: If archaeologists and textual historians together find that a large number of major cities in antiquity had public libraries, and have insufficient data to confirm there was any major city that lacked a public library, then all major cities in antiquity had public libraries.

Minor Premise 2: Archaeologists and textual historians together have found that a large number of major cities in antiquity had public libraries, and have insufficient data to confirm any major city lacked a public library.

Conclusion (= MjP 1): Therefore, all major cities in antiquity had public libraries.

Nesting Test for Major Premise 2:

Minor Premise 3: If a large number of representatives of a class have property x , and there is insufficient data to confirm any members of that class lack x , then all members of that class have x .

Major Premise 3: If it is the case that {MnP 3: if a large number of representatives of a class have property x , and there is insufficient data to confirm any members of that class lack x , then all members of that class have x }, then it is the case that {MjP 2: if archaeologists and textual historians together find that a large number of major cities in antiquity had public libraries, and have insufficient data to confirm there was any major city that lacked a public library, then all major cities in antiquity had public libraries}.

Conclusion (= MjP 2): Therefore, if archaeologists and textual historians together find that a large number of major cities in antiquity had public libraries, and have insufficient data to confirm there was any major city that lacked a public library, then all major cities in antiquity had public libraries.

Nesting Test for Minor Premise 3:

Major Premise 4: If there can be no exceptions to a rule {if A, then B} then it is always the case that {if A, then B}.

Minor Premise 4: There can be no exceptions to the rule {if a large number of representatives of a class have property x , and there is insufficient data to confirm any members of that class lack x , then all members of that class have x }.

Conclusion (= MnP 3): Therefore, if a large number of representatives of a class have property x , and there is insufficient data to confirm any members of that class lack x , then all members of that class have x .

Defeat of Minor Premise 4 (hence of Major Premise 1):

There can be exceptions to the rule {if a large number of representatives of a class have property x , and there is insufficient data to confirm any members of that class lack x , then all members of that class have x }.

Correction of MnP 4 (hence of MjP 1):

There can be no exceptions to the rule {if a large number of representatives of a class have property x , and those members were effectively selected at random from among all members of that class, and there is insufficient data to confirm any members of that class lack x , even though it is somewhat probable we would have such data in at least one instance if many members of that class lacked x , then it is at least somewhat probable that any given member of that class has x }

RESULT: Through nesting we locate the underlying rule, discover its flaw, and when we correct that flaw, we discover the necessary qualifications and analyses we may have overlooked before. For instance, in this case: (a) we now know we should qualify our premises (and thus our conclusion) as a matter of probability, and a probability less than what we would consider a historical certainty; (b) we now know to ask whether we should even *expect* evidence of major cities lacking public libraries, if any did in fact lack them; and (c) we now know to ask: is the sample of major cities for which we have confirmed public libraries effectively a random sample of all major cities—and if not, will the known bias in sampling affect our conclusion?

As to (b), we might be able to say that if there were *many* such deprived cities, we should have evidence of at least *one* case, whereas if there were only a few, we might not have evidence of that, and as to (c) we might observe that the bias now is in fact against having evidence for the largest of cities (since modern cities often stand on top of the most successful ancient cities, making archaeological surveys spotty at best), and since it

is highly improbable that numerous lesser cities would have public libraries while yet greater cities lacked them, the bias is actually *in favor* of our conclusion that all major cities had public libraries.

Therefore, logical analysis can be a useful tool in history: to identify or check against possible errors, by identifying underlying assumptions regarding rules of inference and trends and generalizations in the evidence, and to discover new avenues of analysis, qualification, and inquiry that could improve our methods or results.

5. Syllogistic Representation of Common Historicity Criteria

EXAMPLE 1: The Criterion of Dissimilarity : “If a saying attributed to Jesus is dissimilar to the views of Judaism and to the views of the early church, then it can confidently be ascribed to the historical Jesus”

Major Premise: If any saying x attributed to Jesus is dissimilar to the views of Judaism and to the views of the early church, then Jesus said x .

Minor Premise: Saying x [= Jesus addressing God as his Father] is dissimilar to the views of Judaism and to the views of the early church.

Conclusion: Therefore, Jesus said x [= Jesus addressed God as his Father].

Nesting Test for the Minor Premise:

Major Premise 2: If we have no evidence of saying x [addressing God as one’s Father] from Jews (prior to or contemporary with Jesus) or the early Church (without attribution to Jesus), then saying x [Jesus addressing God as his Father] is dissimilar to the views of Judaism and to the views of the early church.

Minor Premise 2: We have no evidence of saying x [addressing God as one’s Father] from Jews (prior to or contemporary with Jesus) or the early Church (without attribution to Jesus)

Conclusion (MnP): Therefore, Saying x [Jesus addressing God as his Father] is dissimilar to the views of Judaism and to the views of the early church.

Etc.

RESULT: If we continue this nesting analysis we will find both MjP 2 and MnP 2 insupportable: because we *do* have evidence of Jews addressing God as one’s Father, and it is *not* the case that if we have no evidence of a practice that it did not exist.¹ The Criterion of Dissimilarity thus reduces to an *Argumentum ad Ignorantiam*, a textbook fallacy. Applying the criterion to produce a conclusion of any confidence requires just as much confidence that the practice didn’t exist, which is very difficult to establish (one

¹ Cf. Mary Rose D’Angelo, “*Abba* and Father: Imperial Theology in the Contexts of Jesus and the Gospels,” *The Historical Jesus in Context*, edited by Amy-Jill Levine, Dale C. Allison, Jr., and John Dominic Crossan (Princeton University Press: 2006): pp. 64-78.

must thoroughly survey all relevant evidence and scholarship, no simple task) and often *impossible* to establish (since we know for a fact there was a great deal more diversity in Jewish beliefs and practice than we presently know any specifics of, and since the survival of sources is so spotty, no valid conclusion can be reached about what *no* Jews ever thought, said, or did).

That invalidates the Minor Premise (“Saying *x* [Jesus addressing God as his Father] is dissimilar to the views of Judaism and to the views of the early church.”). But even the Major Premise here will be found indefensible on a complete nesting analysis. “If any saying *x* attributed to Jesus is dissimilar to the views of Judaism and to the views of the early church, then Jesus said *x*” assumes an invalid rule of inference: that only Jesus could innovate. But if Jesus could innovate a saying, then so could anyone, including an actual Gospel author (or other intermediary source)—note how Paul innovated a law-free Gentile mission, and if Paul could do that, so could anyone innovate anything, and we know too little about the many Christians and Jews who lived prior to the Gospels to rule any of them out, yet we would have to rule them all out in order to isolate Jesus as the only available innovator we can credit for the innovation.

For instance, any saying (*x*) we think we can isolate as being unique to Jesus, may in fact be unique to Peter instead, who uniquely imagined or hallucinated or dreamed or invented Jesus saying (*x*). There is no more reason to assume the innovation was of Jesus’ own making than of Peter’s (whether consciously, for a specific innovative purpose, or unconsciously, as a construct of what Peter took to be visions or revelations, but were actually the product of his subconscious creatively responding to the problems and ambitions of his time). So how are we to tell the difference?

EXAMPLE 2: Multiple Attestation : “If a tradition is independently attested in more than one source, then it is more likely to be authentic than a tradition attested only once.”

Major Premise: If a tradition is independently attested in more than one source, then it is more likely to be authentic than a tradition attested only once.

Minor Premise: Jesus healing the sick is independently attested in more than one source.

Conclusion: Therefore, Jesus healed the sick.

RESULT: A nesting analysis will find flaws in both the major and minor premise here.

As to the Minor Premise: Jesus healing the sick appears only in the Gospels and later documents influenced by the Gospels, and not in any of the prior Epistles. The Synoptic and Apocryphal Gospels all show mutual dependence and therefore are not independent sources, and modern scholarship has established that the Gospel of John was probably influenced by that of Mark and Luke, and therefore there are no independent attestations of Jesus healing the sick: that concept appears only in documents ultimately derivative of the Gospel of Mark, which is the earliest mention of the tradition. Therefore the criterion of multiple attestation might not apply.

As to the Major Premise: it is well-documented that even a false claim can be multiply attested in independent sources (e.g. multiple independent sources attest to the labors of Hercules), and folklorists have documented that this can occur very rapidly (there is no actual limit to how rapidly multiple sources can transform and transmit the same story). Therefore the rule of inference this criterion depends on is invalid. A more

rigorous rule must be developed that can distinguish between multiple attestation that is caused by the authenticity of what is reported and multiple attestation that is caused by spreading stories of false origin. Such a rule might not be possible, or will be very limited in application because of the extent of the qualifications and conditions that must be met.

A third flaw is that “more likely to be authentic” is vague as to whether (or when) “more likely” means “likely enough to warrant believing it’s true.” When is the evidence enough to warrant belief? None of the historicity criteria developed provide any help in answering this question. But Bayes’ Theorem does.

EXAMPLE 3: The Criterion of Embarrassment : “Since Christian authors would not invent anything that would embarrass them, anything embarrassing in the tradition must be true.”

Major Premise 1: Christians would not invent anything that would embarrass them.

Minor Premise 1: The crucifixion of Jesus would embarrass Christians.

Conclusion 1: Therefore, Christians did not invent the crucifixion of Jesus.

Major Premise 2: A report is either invented or it is true.

Minor Premise 2 (= Conclusion 1): The crucifixion of Jesus was not invented.

Conclusion 2: Therefore, the crucifixion of Jesus is true.

Another way to test rules of inference is to try them out on contrary cases. For example:

Major Premise 1: Cybeleans would not invent anything that would embarrass them.

Minor Premise 1: The castration of Attis would embarrass Cybeleans.

Conclusion 1: Therefore, Cybeleans did not invent the castration of Attis.

Major Premise 2: A report is either invented or it is true.

Minor Premise 2 (= Conclusion 1): The castration of Attis was not invented.

Conclusion 2: Therefore, the castration of Attis is true.

RESULT: This is obviously not a credible conclusion. We have no good reason to believe there was ever an actual Attis who was castrated and it is commonly assumed the story was invented for some particular symbolic reason. The same, then, could be true of the crucifixion of Jesus. Tacitus reports that the castration of Attis was indeed embarrassing (it is grounds for his disgust at the religion), yet the castration of Attis is not a credible story, therefore the criterion of embarrassment is in some manner fallacious.

An example within the Christian tradition is the astonishing stupidity of the Disciples, especially in the earliest Gospel of Mark. Their depiction is in fact so unrealistic it isn’t credible (real people don’t act like that), which means Mark (or his sources) invented that detail *despite* its potential embarrassment. Hence the flaw in the criterion of embarrassment is in assuming that historical truth is the only factor that can overcome the potential embarrassment of some reported detail, when in fact moral or doctrinal or symbolic truth can *also* override such concerns.

For example, Dennis MacDonald argues this attribute emulates the equally-unrealistic stupidity of the crew of Odysseus and thus stands as a marker of the same

things that *their* stupidity represented. That may be true. But I also argue it furthers a literary theme found throughout Mark of the Reversal of Expectation.² Thus everything that seems embarrassing in Mark might be an intentional fabrication meant to convey a lesson. Mark echoes the gospel theme that “the least shall be first” in his construction of all his stories: although Jesus tells Simon Peter he must take up the cross and follow him, Simon the Cyrenean does this instead; although the pillars James and John debate who will sit at Jesus’ right and left at the end, instead two nameless thieves sit at his right and left at the end; although the lofty male Disciples flee and abandon Jesus, the lowly female followers remain faithful, and as a result the *least* are the *first* to discover that Christ is risen; and while Mark begins his Gospel with the “good news” of the “voice crying out” of the lone man who boldly came forward as a “messenger who will prepare our way,” he ends his Gospel with *several women, fleeing in fear and silence, and not* delivering the good news, exactly the opposite of how his book began. So since details that seem embarrassing in Mark might serve his literary intentions, we can’t be certain they’re true.

This final example exposes the importance of testing criteria by comparing them with alternative theories of the evidence. You must ask yourself, *what if I’m wrong?* What *other* reasons might Christians have for inventing potentially embarrassing stories? And how do those reasons compare with the theory that they reported embarrassing stories because they were true? Bayes’ Theorem suits exactly such an analysis.

6. Recommended Reading on Bayes’ Theorem

Eliezer Yudkowsky, “An Intuitive Explanation of Bayesian Reasoning (Bayes’ Theorem for the Curious and Bewildered: An Excruciatingly Gentle Introduction)” at yudkowsky.net/bayes/bayes.html

Douglas Hunter, *Political [and] Military Applications of Bayesian Analysis: Methodological Issues* (Westview Press: 1984).

Luc Bovens and Stephan Hartmann, *Bayesian Epistemology* (Oxford University Press: 2003).

Timothy McGrew, “Bayesian Reasoning: An Annotated Bibliography” at homepages.wmich.edu/~mcgrew/bayes.htm

Wikipedia on “Bayes’ Theorem” (for English: en.wikipedia.org/wiki/Bayes'_theorem)

² See my contributions to *The Empty Tomb: Jesus beyond the Grave*, edited by Jeff Lowder and Robert Price (Prometheus: 2005) and my online book *Was Christianity Too Improbable to be False?* (The Secular Web: 2006) at www.infidels.org/library/modern/richard_carrier/improbable.

7. Bayes' Theorem (Complete)

$$P(h|e.b) = \frac{P(h|b) \times P(e|h.b)}{[P(h|b) \times P(e|h.b)] + [P(\sim h|b) \times P(e|\sim h.b)]}$$

8. Explanation of the Terms in Bayes' Theorem

P = Probability (more specifically in this case, epistemic probability, i.e. the probability that something stated is true)

h = hypothesis being tested

~h = all other hypotheses that could explain the same evidence (if *h* is false)

e = all the evidence directly relevant to the truth of *h* (*e* includes both what *is* observed *and* what is *not* observed, despite ample looking)

b = total background knowledge (all available personal and human knowledge about anything and everything, from physics to history)

P(h|e.b) = the probability that a hypothesis (*h*) is true given all the available evidence (*e*) and all our background knowledge (*b*)

P(h|b) = **the prior probability that *h* is true** = the probability that our hypothesis would be true given only our background knowledge (i.e. if we knew nothing about *e*)

P(e|h.b) = **the posterior probability of the evidence (given *h* and *b*)** = the probability that all the specific evidence we have would exist (or that something comparable to it would exist) if the hypothesis is true, and all our background knowledge is true (i.e. everything else we know about the world, about people, about the time and place in question, etc.) = [hereafter I call this the consequent probability]

P(~h|b) = 1 - P(h|b) = **the prior probability that *h* is false** = the sum of the prior probabilities of all alternative explanations of the same evidence (e.g. if there is only one viable alternative, this means the prior probability of all other theories is vanishingly small, i.e. substantially less than 1%, so that P(~h|b) is the prior probability of the one viable competing hypothesis; if there are many viable competing hypotheses, they can be subsumed under one group category (~*h*), or treated independently by expanding the equation, e.g. for three competing hypotheses [P(h₁|b) x P(e|h₁.b)] + [P(h₂|b) x P(e|h₂.b)] + [P(h₃|b) x P(e|h₃.b)])

P(e|~h.b) = **the posterior probability of the evidence if *b* is true but *h* is false** = the probability that all the specific evidence we have would exist (or that something

Since mathematicians use posterior probability to refer to the final probability (the probability of *h* posterior to *e*), to avoid confusion I will hereafter use consequent probability to refer to what I here describe as the posterior probability of the evidence (the probability of *e* posterior to *h*).

comparable to it would exist) if the hypothesis we are testing *is false*, but all our background knowledge is still true (i.e. everything else we know about the world, about people, about the time and place in question, etc.). This equals the posterior probability of the evidence if some hypothesis *other* than *h* is true, and if there is more than one viable contender, you can include each competing hypothesis independently (per above) or subsume them all under one group category ($\sim h$).
 = [hereafter I call this the consequent probability]

9. Bayes' Theorem (Abbreviated)

$$P(h|e.b) = \frac{P(h|b) \times P(e|h.b)}{P(e|b)}$$

Note that $P(e|b) = [P(h|b) \times P(e|h.b)] + [P(\sim h|b) \times P(e|\sim h.b)]$ so I recommend the longer form instead, as a check against error (especially that of assuming $P(e|b) = P(e|\sim h.b)$).

10. Employing Bayes' Theorem

If we use Bayes' Theorem to determine the likelihood that Jerusalem had a public library, and if the following data is the same (note that this model of the problem and these percentages are deliberately unrealistic—realistically they should all be much higher, and the evidence is more complicated):

- 60% chance that any major city in antiquity had a public library
- 90% chance that Jerusalem was a major city
- 60% chance that there was a library if there is a document attesting to a library
- 90% chance that we have a document attesting to a public library in Jerusalem

$P(h|b) = (0.6 \times 0.9) + x = 0.54 + x$ = The prior probability that Jerusalem was a major city and (as such) would have a public library (based on our background knowledge alone, i.e. what we know of other cities and of the archaeological and textual evidence of the size and success of Jerusalem as a city). The added factor x is the prior probability that Jerusalem had a public library even if it *wasn't* a major city. Here for simplicity this is arbitrarily assumed to be zero, but realistically it needn't be. And if it wasn't, x would equal $0.1 \times y$, where 0.1 is the probability that Jerusalem *wasn't* a major city (i.e. the converse of the probability that it *was*, i.e. 0.9) and y is the probability of a non-major city having a public library (whatever we deemed that to be from available evidence).

$P(\sim h|b) = 1 - 0.54 = 0.46$ = The prior probability that Jerusalem did *not* have a public library = the converse of the other prior (i.e. all prior probabilities that appear in a Bayesian equation must sum to exactly 1, no more nor less, because together they must encompass all possible explanations of the evidence for Bayes' Theorem to be valid).

$P(e|h.b) = 1.0$ = The consequent probability that we would have either some specific evidence of a public library at Jerusalem *and / or* no evidence against there being one. For example, if we assume there is an 80% chance that no evidence would survive even if there was a library there, then there is a 20% chance that some evidence would survive, but this doesn't mean the survival of such evidence drops $P(e|h.b)$ to 20% (as if having no evidence made the library more likely than if we had evidence). For on h we can expect *either* no evidence *or* some evidence, so the total probability of having what we have (either none or some) is 100% (since having some constitutes having 'either none or some'). Notably this renders irrelevant the question of whether we actually do have a document attesting to a library, so the probability that we do makes no difference to the likelihood of h . The only thing we *don't* expect on h is evidence *against* there being a public library at Jerusalem, any of which (no matter how indirect or indefinite) would reduce $P(e|h.b)$ below 100% (to whatever degree was appropriate, i.e. in accordance with how strongly or surely or clearly that evidence argues against a public library there).

$P(e|\sim h.b) = 0.4$ = The consequent probability that there would be a document attesting to a public library in Jerusalem even when there wasn't a public library there (i.e. the converse of the 60% chance that such a library existed if we have such a document).

The probability that we don't have such a document can affect $P(e|\sim h.b)$, by upping it slightly. But that would require a longer digression.

$$\begin{aligned}
 P(h|e.b) &= \frac{P(h|b) \times P(e|h.b)}{[P(h|b) \times P(e|h.b)] + [P(\sim h|b) \times P(e|\sim h.b)]} \\
 P(h|e.b) &= \frac{0.54 \times 1.00}{[0.54 \times 1.00] + [0.46 \times 0.40]} = \frac{0.54}{[0.54] + [0.184]} \\
 &= \frac{0.54}{[0.724]} = 0.746 = 75\% = P(h|e.b)
 \end{aligned}$$

RESULT: Instead of the implausible 29% chance that Jerusalem had a public library, we get the plausible result of a 75% such chance. And this was found with unrealistic percentages that biased the result *against* there being a library there, which means *a fortiori* we can be highly certain Jerusalem had a public library. This illustrates the advantage of using unrealistically hostile estimates against a hypothesis, since if the conclusion follows even then, we can have a high confidence in that conclusion.

11. The Advantages of Bayes' Theorem

1. Bayes' Theorem will help you determine how to tell if your theory is *probably* true rather than merely *possibly* true.

It achieves this (a) by forcing you to compare the relative likelihood of different theories of the same evidence (so you *must* think of *other* reasons the evidence we have might exist, *besides* the reason you intuitively think is most likely), and (b) by forcing you to examine what exactly you mean when you say something is 'likely' or 'unlikely' or is 'more likely' or 'less likely'. With Bayes' Theorem you have to think in terms of relative probabilities, as in fact must be done in all sound historical reasoning, which ultimately requires matching numbers (or ranges of numbers) to your vocabulary of likelihood.

2. Bayes' Theorem will inspire a closer examination of your background knowledge, and of the corresponding objectivity of your estimates of prior probability.

Whether you are aware of it or not, all your thinking relies on estimations of prior probability. Making these estimations explicit will expose them to closer examination and test. Whenever you say some claim is implausible or unlikely because 'that's not how things were done then', or 'that's not how people would likely behave', or 'other things happened more often instead', you are making estimates of the prior probability of what is being claimed. And when you make this reasoning explicit, unexpected discoveries can be made.

For example, as Porter and Thiessen have both observed, it's inherently *unlikely* that any Christian author would include *anything* embarrassing in a written account of his beliefs, since he could choose to include or omit whatever he wanted. In contrast, it's inherently *likely* that anything a Christian author included in his account, he did so for a deliberate reason, to accomplish something he *wanted* to accomplish, since that's how all authors behave, especially those with a specific aim of persuasion or communication of approved views. Therefore, already the prior probability that a seemingly embarrassing detail in a Christian text is in there because it is true *is low*, whereas the prior probability that it is in there for a specific reason *regardless* of its truth *is high*.

3. Bayes' Theorem will force you to examine the likelihood of the evidence on competing theories, rather than only one—in other words, forcing you to consider what the evidence should look like if your theory happens to be false (What evidence can you then expect there to be? How would the evidence in fact be different?). Many common logical errors are thus avoided. You may realize the evidence is just as likely on some alternative theory, or that the likelihood in either case is not sufficiently different to justify a secure conclusion.

For example, Paul refers to James the Pillar as the Brother of the Lord, and to the Brothers of the Lord as a general category of authority besides the Apostles. It is assumed this confirms the historicity of Jesus. But which is more likely, that a historical (hence biological) brother of Jesus would be called the Brother of the Lord, or that he would be called the Brother of *Jesus*? In contrast, if we theorize that 'Brother of the Lord' is a rank in the Church, not a biological status, then the probability that we would hear of

authorities being called by that title is just as high, and therefore that Paul mentions this title is not by itself sufficient evidence to decide between the two competing theories of how that title came about.

Estimates of prior probability might then decide the matter, but one then must undertake a total theory of the evidence (extending beyond just this one issue), since there is no direct evidence here as to what was normal (since there is no precedent for calling anyone “Brother of the Lord” as a biological category, and only slim or inexact precedent for constructing such a title as a rank within a religious order).

4. Bayes’ Theorem eliminates the Fallacy of Diminishing Probabilities.

Bayes’ Theorem requires a total theory of the evidence, as all historical reasoning should, rather than focusing on isolated pieces of information without regard for how they all fit together. But it does this by balancing every term in the numerator with a term in the denominator.

For example, in the public libraries argument we saw that adding two pieces of evidence together reduced the probability of the conclusion, contrary to common sense. It would have been even worse if we had *ten* items of evidence. But in Bayes’ Theorem, for every element of P(e|h.b) there is a corresponding element of P(e|~h.b), producing the mathematical result that the more evidence you add, the *higher* the probability of the hypothesis, exactly as we should expect.

For instance, if we had four items of evidence, each 60% likely if *h* is true, on a straight calculation the probability of having all four items of evidence on *h* is 0.6 x 0.6 x 0.6 x 0.6 = 0.6⁴ = 0.1296 = 13%, which means the more evidence we have, the less likely *h* is, contrary to reason. But on Bayes’ Theorem we ask how likely that same evidence is if *h* is false (and any other hypothesis is true instead). If each item of evidence in our hypothetical case was only 40% likely on ~*h*, then Bayes’ Theorem would give us (for simplicity’s sake assuming the prior probabilities are equal):

$$P(h|e.b) = \frac{P(h|b) \times P(e|h.b)}{[P(h|b) \times P(e|h.b)] + [P(\sim h|b) \times P(e|\sim h.b)]}$$

$$P(h|e.b) = \frac{0.5 \times [0.6 \times 0.6 \times 0.6 \times 0.6]}{[0.5 \times [0.6 \times 0.6 \times 0.6 \times 0.6]] + [0.5 \times [0.4 \times 0.4 \times 0.4 \times 0.4]]}$$

$$P(h|e.b) = \frac{0.5 \times 0.1296}{[0.5 \times 0.1296] + [0.5 \times 0.0256]} = \frac{0.0648}{0.0648 + 0.0128} = \frac{0.0648}{0.0776}$$

$$P(h|e.b) = 0.835 = 84\%$$

RESULT: The more evidence you have, the higher the probability of the hypothesis, exactly as common sense would expect.

5. Bayes' Theorem has been proven formally valid. Any argument that violates a valid form of argument is itself invalid. Therefore any argument that violates Bayes' Theorem is invalid. All valid historical arguments are described by Bayes' Theorem. Therefore any historical argument that cannot be described by a correct application of Bayes' Theorem is invalid. Therefore Bayes' Theorem is a good method of testing any historical argument for validity.

6. You can use Bayesian reasoning without attempting any math (see below), but the math keeps you honest, and it forces you to ask the right questions, to test your assumptions and intuitions, and to actually give relative weights to hypotheses and evidence that are not all equally likely. *But above all, using Bayes' Theorem exposes all our assumptions to examination and criticism, and thus allows progress to be made, as we continually revise our arguments in light of the flaws detected in our reasoning.*

12. Avoiding Common Errors with Bayes' Theorem

1. The Fallacy of False Precision.

One common objection to using Bayes' Theorem in history is that Bayes' is a model of mathematical precision in a field that has nothing of the kind. This precision of the math can create the illusion of precision in the estimates and results. But as long as you do not make this mistake, it will not affect your results.

The correct procedure is to choose values for the terms in the equation that are at the limit of what you can reasonably believe them to be, to reflect a wide margin of error, thus ensuring a high confidence level (producing an argument *a fortiori*), regardless of any inexactness in your estimations.

For example, surely more than 60% of major cities in antiquity had public libraries (the evidence is compellingly in favor of a much higher percentage, provided 'major city' is reasonably defined). But since we don't have exact statistics, we can say that the percentage of such cities must fall between 60% and 100% (= 80% with a margin of error +/-20%). With such a wide margin of error, our confidence level remains high (see appendix). We are in effect saying that we might not be sure it was 100% (or 90% or even 80%), even though we may

believe it is, but we *can* be sure it was *no less* than 60%. Since that is the limit of what we deem reasonable, so will our conclusion be (the conclusion is only as strong as an argument's weakest premise, and each probability assigned in a Bayesian equation is the formal equivalent of a premise).

2. The Fallacy of Confusing Evidence with Theories.

A single example will suffice: William Lane Craig frequently argues that historians need to explain the evidence of the empty tomb. But in a Bayesian equation, the evidence is not the discovery of an empty tomb, but the production of a *story* about the discovery of an empty tomb. That there was an actual empty tomb is only a theory (a hypothesis, i.e. h) to explain the production of the story (which is an element of e). But this theory must be compared with *other* possible explanations of why that story came to exist ($= \sim h$, or $= h_2, h_3$, etc.), and these must be compared on a total examination of the evidence (*all* elements of e , in conjunction with b and the resulting prior probabilities).

Hence a common mistake is to confuse actual hypotheses about the evidence, with the actual evidence itself (which should be tangible physical facts, i.e. actual surviving artifacts, documents, etc., and straightforward generalizations therefrom). Though hypotheses can in principle be treated as evidence, this is often only mathematically practical (or non-fallacious) when such hypotheses are so well confirmed as to be nearly as certain as the existence of the evidence that confirms them.

3. The Fallacy of Confusing Assumptions with Knowledge.

Assumptions in, assumptions out. Therefore, assumptions usually should have no place in Bayesian argument, as its conclusions will only be as strong as their weakest premise, and an assumption is a very weak premise indeed.

The term b in Bayes' Theorem establishes that all the probabilities in the equation are conditional probabilities, i.e. probabilities conditional on the truth of our background knowledge. Therefore, only background *knowledge* should be included in b and thus considered in assigning probabilities, not background assumptions or dogmatic beliefs. The difference between professional and unprofessional history is the acceptance in b of only what has been more or less accepted by peer review as an established fact (although the possibility of something can be accepted even when the certainty of that something is not, but in any case the only position that counts as professional when determining background knowledge is the position all experts can agree is acceptable).

This error leads to a common misapprehension that, for example, prior probabilities in Bayes' Theorem are worldview-dependent. They are not. For example, it doesn't matter whether you are a naturalist and believe no miracles exist, or a Christian and believe they do. Either way, if you are acting professionally, you both must agree that so far as is *objectively known*, most miracle claims in history have turned out to be bogus and none have been confirmed as genuine, therefore the prior probability that a miracle claim is

genuine must reflect the *fact* that most miracle claims are not, and that is a fact even if genuine miracles exist.

In other words, the naturalist must allow that he could be wrong (so he must grant *some* probability that there might still be a genuine miracle somewhere, whatever that probability must be) and the Christian must allow that most miracle claims are false (not only because investigated claims trend that way, but also because the Christian already grants that most miracle claims validate other religious traditions, and therefore must be false if Christianity is true, and many even within the Christian tradition strain even a Christian's credulity, and therefore are probably false even if Christianity is true). If most miracle claims are false, then the prior probability that any miracle claim is false must be high, regardless of whether miracles exist or whether Christianity is true.

So although worldview considerations *can* be brought into *b*, Bayes' Theorem does not require this. And when such considerations *are* brought into *b*, that only produces conditional probabilities that follow only when the adopted worldview is true. But if a certain worldview is already assumed to be true, most arguments don't even have to be made, as the conclusion is already foregone. Therefore, *b* should include only objectively agreed knowledge (and probabilities assessed accordingly), unless arguing solely to audiences *within* a single worldview community.

NOTE: Further discussion of all these topics (and in considerable detail) will appear in my forthcoming book: Richard Carrier, *On the Historicity of Jesus Christ* (there especially in chapter two and a mathematical appendix). For contact information and more about my work in other aspects of Jesus studies and beyond, see www.richardcarrier.info.

Additional Notes for Further Consideration:

A. Bayesian Reasoning without Mathematics

1. You don't have to use scary math to think like a Bayesian, unless the problem is highly complex, or you want clearer ideas of relative likelihood. For instance...
2. Easy Case of Rejection: If you estimate that the prior probability of *h* must be at least somewhat low (any degree of "*h* is unlikely, given just *b*"), and you estimate the evidence is no more likely on *h* than on $\sim h$, then *h* is probably false ("*h* is unlikely even given *e*").

EXAMPLE: Talking donkeys are unlikely, given everything we know about the world. That there would be a story about a talking donkey is just as likely if there were a real talking donkey than if someone just made up a story about a talking donkey. Therefore, *e* is no more likely on *h* (Balaam's donkey actually

spoke) than on $\sim h$ (Someone made up a story about Balaam's donkey speaking), and h (Balaam's donkey actually spoke) is already initially unlikely, whereas $\sim h$ (Someone made up a story about Balaam's donkey speaking) is initially quite likely (since on b we know people make up stories all the time, but we don't know of any talking donkeys). Therefore, we can be reasonably certain that Balaam's donkey didn't talk. Note how this conclusion is worldview-independent. It follows from plain facts everyone can agree on.

3. Unexpected Case of Acceptance: If you estimate that the prior probability of h must be at least somewhat high (any degree of " h is likely given just b "), even if you estimate the evidence is no more likely on h than on $\sim h$, then h is still probably true (" h is likely even given e ").

EXAMPLE: That Julius Caesar regularly shaved (at least once a week) is likely, given everything we know about Julius Caesar and Roman aristocratic customs of the time and human male physiology. That we would have no report of his shaving anytime during the week before he was assassinated is as likely on h (Caesar shaved during the week before he was assassinated) as on $\sim h$ (Caesar didn't shave any time during the week before he was assassinated), since, either way, we have no particular reason to expect any report about this to survive. Nevertheless, h is probably true: Caesar shaved sometime during the week before he was assassinated.

4. In the above examples, exact numbers and equations were unneeded, just the innate logic of Bayesian reasoning sufficed. This is the case for most historical problems. Only when the problem is complex does math become a necessary tool.

For example, if you want some idea of *how likely* a hypothesis is, then you may need to do some math, unless the relative degrees of likelihood are so clear you can reason out the result without any equation. For instance, it is *very unlikely* that any donkey spoke, therefore it is *very unlikely* that Balaam's story is true. But things are not always this clear.

Or, for example, when the prior probability of a hypothesis is low, but the evidence is still much more likely on that hypothesis than any other (or vice versa), then you need to evaluate how low and how high these probabilities are, in order to determine if they balance each other out.

A flow chart results, wherein the mathematical terms (as given below) can be replaced with ordinary sentences in your native language:

i.e. For a high $P(h|e.b)$, $P(e|\sim h.b)$ must be a lot lower than $P(e|h.b)$, in fact the lower $P(h.b)$ is, the lower $P(e|\sim h.b)$ must then be.

i.e. For a high $P(h|e.b)$, $P(e|\sim h.b)$ must not be higher than $P(e|h.b)$ by too much, but the higher $P(h.b)$ is, the higher $P(e|\sim h.b)$ can be and still not be too much.

